IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant(s):	Hari K. Ravichandran		
Assignee:	Sun Microsystems, Inc.		
Title:	Apparatus and Method for Processor Performance Monitoring		
Serial No.:	10/056,294	Filing Date:	January 22, 2002
Examiner:	Aaron D. Matthew	Group Art Unit:	2114
Docket No.:	P2678	Customer No.:	33438

Austin, Texas
August 1, 2005

MAIL STOP AF
COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, VA 22313-1450

PRE-APPEAL BRIEF REQUEST FOR REVIEW AND ARGUMENTS

Sir:

This paper accompanies the Pre-Appeal Brief Request for Review and sets forth a succinct, concise, and focused set of arguments for which the review is being requested.

CLAIM STATUS

Claims 9-19 are pending.

Claims 9 stands rejected under Bunnell, U.S. Patent No. 5,564,015 (Bunnell) in view of Roeber, et al., U.S. Patent No. 5,682,328 (Roeber). Claim 13 stands rejected under Bunnell, in view of Roeber, in further view of Levine et al., U.S. Patent No. 6,067,644 (Levine).

ARGUMENTS

Independent claim 9 is allowable over Bunnell and Roeber.

Bunnell discloses a central processing unit ("CPU") activity monitor and method providing CPU activity information. The CPU activity monitor includes a timer and an activity event counter for receiving a plurality of mode signals from the CPU, a cache miss signal from a cache memory system, and a clock signal from a clock. An activity-to-inactivity value defines

when the CPU transitions from an active state to an inactive state. An activity threshold defines when the CPU transitions from an inactive state to an active state.

More specifically, the examiner has stated that Bunnell discloses generating a first high speed memory miss signal, a first high speed memory miss count signal and a time stamp signal. The examiner cites the following portions of Bunnell:

> Thus, the cache memory system 30 generates the cache miss signal 36 when the CPU 28 accesses data that is not stored in the cache memory system 30. (Bunnell, Col. 6, lines 58 – 61.)
>
> The present invention determines CPU activity by counting the number of activity events that occur within a certain time interval. The preferred embodiment defines a CPU activity event as a CPU data write cycle that generates a cache miss signal. (Bunnell, Col. 4. lines 49 – 53.)
>
> The output of the event register 112 and the output of the event counter 114 connect to the event comparator 104. (Bunnell, Col. 9, lines 5 – 6.)
>
> In addition, the CPU activity monitor receives a clock signal. (Bunnell, Col. 4, lines 44-45.)

Roeber, et al. discloses event logging using a single board computer control card configurable onto a backplane containing target processors being monitored. A high resolution clock on the control card time stamps events. Memory on the control card provides a central buffer to store event data and stores a control program effecting functionality of the card. A network interface facilitates communication with host computers for post processing of event data and to control, communicate with, and access the control card. A control program effects event data collection and organization/storage of events in control card memory. The control program coordinates retrieval of events from an event interface area of memory on slave target processors prior to processing by the control card. The control program coordinates offloading of event data from the control card to host computers for post processing by known software visualization tools. Target software is instrumented with calls to a macro which in turn calls the logging function to effect recording of events. The event records are temporarily stored in an event interface portion of memory on the control card or in a buffer on the target processors.

More specifically, the examiner has stated that Roeber discloses a method for monitoring and analyzing system activity by recording event data along with time information associated with the event data. The examiner cites the following portion of Roeber:

A high resolution clock on the control card is used to time stamp events. A portion of memory on the control card is used as a central buffer to store event data, while another portion of memory on the control card is used to store a control program that effects functionality of the control card. (Roeber, Col. 3, lines 30 – 33.)

The record created will include the time that the event occurred based on the system clock, and may also include some data relevant to the event. (Roeber, Col. 1, lines 33-35.)

The examiner thus maintains that:

One of ordinary skill in the art would have been motivated to include a second cache in the method disclosed in Bunnell, in view of Roeber et al, in order to improve system performance. Moreover, one of ordinary skill in the art would have considered it obvious to perform the same steps in analyzing cache miss activity associated with the second memory device as were performed in analyzing cache miss activity associated with the first. (Final Office Action, page 9.)

However, Bunnell and Roeber, taken alone or in combination, do not teach or suggest a method for monitoring an execution of a program which includes *when the entry in the first memory device does not exist, generating at least one probe signal indicating a miss entry in the first memory device*; generating a temporal identifier signal that is ***associated with the probe signals***; and ***storing the temporal identifier signal and the probe signals*** *in memory*, all as required by claim 9. Accordingly, claim 9 is allowable over Bunnell and Roeber. Claims 10 – 12 depend from claim 9 and are allowable for at least this reason.

Independent claim 13 is allowable over Bunnell, Roeber and Levine.

Bunnell and Roeber are discussed above.

Levine discloses a processor operable for processing an instruction through a plurality of internal stages will produce a result of the processing of the process at each stage or a reason code why the stage was unable to process the instruction. The result or the reason code will then be passed to a subsequent stage, which will attempt to process the instruction. The second stage will forward the reason code when it cannot produce its own result and it is idle. The second stage will create its own reason code when it is not idle but cannot produce a result, and will forward this reason code.

The examiner has stated that Levine teaches a method of monitoring the execution of instructions in a program including the steps of checking a second cache in the event that an

entry in the first cache dies not exist. The examiner cites to the following portion of Levine:

> Most computer memory systems take advantage of this fact by incorporating small staging areas. These staging areas store frequently used data. Such areas are usually smaller and more rapidly accessed than system memory. This allows systems to complete work faster and thus have higher performance. Similarly, some items held in the staging areas are more frequently used than others. This leads to the use of an additional secondary staging area. If the required datum is not in the first staging area, the second staging area is next checked. If the item is not in the second staging area, system memory is checked. Because of the high probability of finding a required datum in some staging area, the average time to retrieve the datum is typically lower in hierarchically configured memory systems. Consequently, current memory systems are structured as hierarchies of staging areas where the staging areas become larger and slower in order of access.
>
> It is clearly desirable to keep the most frequently reused data in the staging areas (hereinafter referred to as "caches") closets to point of usage. (Levine, Col. 1, line 51 – Col. 2, line 3.)

However, it does not follow that it would have been obvious to analyze cache miss activity for multilevel or multiple caches as claimed merely because Levine discloses two levels of caches.

More specifically, Bunnell, Roeber and Levine, taken alone or in combination, do not teach or suggest a method for monitoring an execution behavior of a program, which includes *generating a second high-speed memory miss signal,* ***a second high-speed memory miss count signal*** *and a time stamp signal, the second high-speed memory miss signal indicating a miss in a second high-speed memory,* ***the second high-speed memory miss count signal representing a number of misses in the second high-speed memory,*** *and* ***the time stamp signal indicating when the second high-speed memory miss signal is active***, all as required by claim 13. Accordingly, claim 13 is allowable over Bunnell, Roeber and Levine. Claims 14 - 19 depend from claim 13 and are allowable for at least this reason.

In light of the above remarks, Applicants respectfully request withdrawal of the rejections.

I hereby certify that this correspondence is being deposited with the United States Postal Service as First Class Mail in an envelope addressed to: Mail Stop AF, COMMISSIONER FOR PATENTS, P.O. Box 1450, Alexandria, VA 22313-1450, on August 1, 2005.

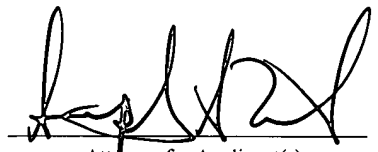

Attorney for Applicant(s)

8/1/05

Date of Signature

Respectfully submitted,

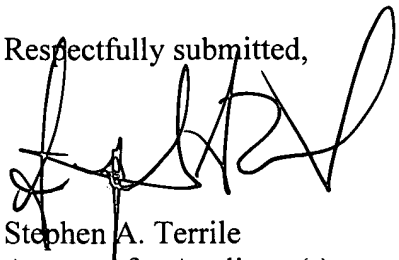

Stephen A. Terrile
Attorney for Applicant(s)
Reg. No. 32,946